

Embedding into the rectilinear plane in optimal $O(n^2)$ time

NICOLAS CATUSSE, VICTOR CHEPOI, and YANN VAXÈS

Laboratoire d'Informatique Fondamentale,
 Université d'Aix-Marseille,
 Faculté des Sciences de Luminy,
 F-13288 Marseille Cedex 9, France
 {catusse,chepoi,vaxes}@lif.univ-mrs.fr

Abstract. In this paper, we present an optimal $O(n^2)$ time algorithm for deciding if a metric space (X, d) on n points can be isometrically embedded into the plane endowed with the l_1 -metric. It improves the $O(n^2 \log^2 n)$ time algorithm of J. Edmonds (2008). Together with some ingredients introduced by Edmonds, our algorithm uses the concept of tight span and the injectivity of the l_1 -plane. A different $O(n^2)$ time algorithm was recently proposed by D. Eppstein (2009).

1. INTRODUCTION

Deciding if a finite metric space (X, d) admits an isometric embedding or an embedding with a small distortion into a given geometric space (usually \mathbb{R}^k endowed with some norm-metric) is a classical question in distance geometry which has some applications in theoretical computer science, visualization, and data analysis. The first question can be answered in polynomial time if \mathbb{R}^k is endowed with the Euclidean metric due to classical results of Menger and Schönberg [6]. On the other hand, by a result of Frechet [6], any metric space can be isometrically embedded into some \mathbb{R}^k with the l_∞ -metric. However, it is NP-hard to decide if a metric space isometrically embeds into some \mathbb{R}^k endowed with the l_1 (alias rectilinear or Manhattan) metric [2, 6]. More recently, Edmonds [9] established that it is even NP-hard to decide if a metric space embeds into \mathbb{R}^3 with l_∞ -metric (a similar question for \mathbb{R}^3 with l_1 -metric is still open). In case of \mathbb{R}^2 , both l_1 - and l_∞ -metrics are equivalent because the second metric can be obtained from the first one by a rotation of the plane by 45° and then by a shrink by a factor $\frac{1}{\sqrt{2}}$. The embedding problem for the rectilinear plane was investigated in the papers [3, 12], which ultimately show that a metric space (X, d) embeds into the l_1 -plane if and only if any subspace with at most six points does [3] (a similar result for embedding into the l_1 -grid was obtained in [4]). As a consequence, it is possible to decide in polynomial time if a finite metric space embeds into the l_1 -plane. Edmonds [9] presented an $O(n^2 \log^2 n)$ time algorithm for this problem and very recently we learned that Eppstein [10] described an optimal $O(n^2)$ time algorithm (for earlier algorithmic results, see also [5]). In this note, independently of [10], we describe a simple and optimal algorithm for this embedding problem, which is different from that of [10].

We conclude this introductory section with a few definitions. In the sequel, we will denote by d_1 or by $\|\cdot\|_1$ the l_1 -metric and by d_∞ the l_∞ -metric. A metric space (X, d) is *isometrically embeddable* into a host metric space (Y, d') if there exists a map $\varphi : X \rightarrow Y$ such that

$d'(\varphi(x), \varphi(y)) = d(x, y)$ for all $x, y \in X$. In this case we say that X is a subspace of Y . A *retraction* φ of a metric space (Y, d) is an idempotent nonexpansive mapping of Y into itself, that is, $\varphi^2 = \varphi : Y \rightarrow Y$ with $d(\varphi(x), \varphi(y)) \leq d(x, y)$ for all $x, y \in Y$. The subspace of Y induced by the image of Y under φ is referred to as a *retract* of Y . Let (X, d) be a metric space. The *(closed) ball* and the *sphere* of center x and radius r are the sets $B(x, r) = \{p \in X : d(x, p) \leq r\}$ and $S(x, r) = \{p \in X : d(x, p) = r\}$, respectively. The *interval* between two points x, y of X is the set $I(x, y) = \{z \in X : d(x, y) = d(x, z) + d(z, y)\}$. Any ball of (\mathbb{R}^k, d_∞) is an axis-parallel cube. A subset S of X is *gated* if for every point $x \in X$ there exists a (unique) point $x' \in S$, the *gate* of x in S , such that $x' \in I(x, y)$ for all $y \in S$ [8]. The intersection of gated sets is also gated. A *geodesic* in a metric space is the isometric image of a line segment. A metric space is called *geodesic* (or *Menger-convex*) if any two points are the endpoints of a geodesic.

For a point p of \mathbb{R}^2 , denote by $Q_1(p), \dots, Q_4(p)$ the four quadrants of \mathbb{R}^2 defined by the vertical and horizontal lines passing via the point p and labeled counterclockwise. Any interval $I_1(x, y)$ of the rectilinear plane (\mathbb{R}^2, d_1) is an axis-parallel rectangle which can be reduced to a horizontal or vertical segment. Any ball of (\mathbb{R}^2, d_1) is a lozenge obtained from an axis-parallel square by a rotation by 45° degrees. In the rectilinear plane, any halfplane defined by a vertical or a horizontal line is gated. As a consequence, axis-parallel rectangles, quadrants, and strips of (\mathbb{R}^2, d_1) are gated as intersections of such halfplanes.

2. TIGHT SPANS

A metric space (X, d) is called *hyperconvex* (or *injective*) [1, 11] if any family of closed balls $B(x_i, r_i)$ with centers x_i and radii r_i , $i \in I$, satisfying $d(x_i, x_j) \leq r_i + r_j$ for all $i, j \in I$ has a nonempty intersection, that is, (X, d) is a geodesic space such that the closed balls have the Helly property. Since the closed balls of (\mathbb{R}^k, d_∞) are axis-parallel boxes, the metric spaces (\mathbb{R}^k, d_∞) and (\mathbb{R}^2, d_1) are hyperconvex. It is well known [1] that (X, d) is hyperconvex iff it is an absolute retract, that is, (X, d) is a retract of every metric space into which it embeds isometrically. As shown by Isbell [11] and Dress [7], for every metric space (X, d) there exists the smallest injective space $T(X)$ extending (X, d) , referred to as the *injective hull* [11], or *tight span* [7] of (X, d) . The tight span of a finite metric space (X, d) can be defined as follows. Let $T(X)$ be the set of functions f from X to \mathbb{R} such that

- (1) for any x, y in X , $f(x) + f(y) \geq d(x, y)$, and
- (2) for each x in X , there exists y in X such that $f(x) + f(y) = d(x, y)$.

One can interpret $f(x)$ as the distance from f to x . Then (1) is just the triangle inequality. Taking $x = y$ in (1), we infer that $f(x) \geq 0$ for all $x \in X$. The requirement (2) states that $T(X)$ is minimal, in the sense that no value $f(x)$ can be reduced without violating the triangle inequality. We can endow $T(X)$ with the l_∞ -distance: given two functions f and g in $T(X)$, define $\rho(f, g) = \max |f(x) - g(x)|$. The resulting metric space $(T(X), \rho)$ is injective and $(T(X), \rho)$ is called the *tight span* of (X, d) . There is an isometric embedding of X into its tight span $T(X)$. Moreover, *any isometric embedding of (X, d) into an injective metric space (Y, d') can be extended to an isometric embedding of $(T(X), \rho)$ into (Y, d')* , i.e., $(T(X), \rho)$ is the smallest injective space into which (X, d) embeds isometrically.

In general, tight spans are hard to visualize. Nevertheless, if $|X| \leq 5$, Dress [7] completely described $T(X)$ via the interpoint-distances of X . For example, if $|X| = 3$, say $X = \{x, y, z\}$, then $T(X)$ consists of three line segments joined at a (Steiner) point, with the points of X




FIGURE 1. Tight span of 3- and 4-point metric space.




FIGURE 2. The three canonical types of tight span of 5-point metric space.

at the ends of the arms (see Fig. 1 (a)). The lengths of these segments are $\alpha_x, \alpha_y, \alpha_z$, where $\alpha_x := (y, z)_x = 1/2(d(x, y) + d(x, z) - d(y, z))$ is the Gromov product of x with the couple y, z (α_y and α_z are defined in a similar way). Notice that one of the values $\alpha_x, \alpha_y, \alpha_z$ may be 0, in this case one point is located between two others. If $|X| = 4$, then the generic form of $T(X)$ is a rectangle $R(X)$ endowed with the l_1 -metric, together with a line segment attached by one end to each corner of this rectangle (see Fig. 1 (b)). The four points of X are the outer ends of these segments. The lengths of these segments and the sides of the rectangle can be computed in constant time from the pairwise distances between the points of X ; for exact calculations see [7]. It may happen that $R(X)$ degenerates into a segment or a point. Finally, there are three canonical types of tight spans of 5-point metric spaces precisely described in [7] (see also Fig. 2 for an illustration). Each of them consists of four or five rectangles, five segments, and eventually one rectangular triangle, altogether constituting a 2-dimensional cell complex. All sides of the cells can be computed in constant time as described in [7]. It was also noticed in [7] that if for each quadruplet X' of a finite metric space (X, d) the rectangle $R(X')$ is degenerated, then (X, d) isometrically embeds into a (weighted) tree and its tight span $T(X)$ is a tree-network.

From the construction of tight spans of 3- and 4-point metric spaces immediately follows that any metric space (X, d) with at most 4 points and its tight span $(T(X), \rho)$ can be isometrically embedded into the l_1 -plane as shown in Fig. 1 (b). This is no longer true for metric spaces on 5 points: to embed, some cells of the tight span must be degenerated. If $|X| = 4$ and the rectangle $R(X)$ is non-degenerated, one can easily show that $R(X)$ isometrically embeds into the l_1 -plane only as an axis-parallel rectangle. Therefore, if additionally the four line segments of $T(X)$ are also non-degenerated, then up to a rotation of the plane by 90° , X and $T(X)$ admit exactly two isometric embeddings into the l_1 -plane. If one corner of $R(X)$ is a point of X and the embedding of the rectangle $R(X)$ is fixed, then there exist three types of isometric embeddings of X and $T(X)$ into the rectilinear plane: two segments of $T(X)$ are embedded as axis-parallel

segments and the third one as a segment whose slope has to be determined. Analogously, if two incident corners of $R(X)$ are points of X , the two segments of $T(X)$ are either embedded as axis-parallel segments, or one as a horizontal or vertical segment and another one as segment whose slope has to be determined. Note also that from the combinatorial characterization of finite metric subspaces of the l_1 -plane presented in [3] immediately follows that a tree-metric (X, d) is isometrically embeddable into the l_1 -plane if and only if the tree-network $T(X)$ has at most four leaves. Finally note that since (\mathbb{R}^2, d_1) is injective, by minimality property of tight spans, $T(X)$ is an isometric subspace of the l_1 -plane for any finite subspace X of \mathbb{R}^2 .

3. ALGORITHM AND ITS CORRECTNESS

3.1. Outline of the algorithm. Let (X, d) be a metric space with n points, called *terminals*. Set $X = \{x_1, \dots, x_n\}$. Our algorithm first finds in $O(n^2)$ time a quadruplet P° of X whose tight span contains a nondegenerated rectangle $R(P^\circ)$. If such a quadruplet does not exist, then (X, d) is a tree-metric and $T(X)$ is a tree-network. If this tree-network contains more than four leaves, then (X, d) cannot be isometrically embedded into the l_1 -plane, otherwise such an embedding can be easily derived. Given the required quadruplet P° , we consider any isometric embedding of P° and of its tight span into the l_1 -plane as illustrated in Fig. 4 and partition the remaining points of X into groups depending on their location in the regions of the plane defined by the rectangle $R(P^\circ)$ and the segments of $T(P^\circ)$. The exact location of points of X in these regions is uniquely determined except the four quadrants defined by $R(P^\circ)$. At the second stage, we replace the quadruplet P° by another quadruplet P by picking one furthest from $R(P^\circ)$ point of X in each of these quadrants. We show that the rectangle $R(P^\circ)$ is contained in the rectangle $R(P)$, moreover, for any isometric embedding φ_0 of P and $T(P)$ into the l_1 -plane, the quadrants defined by two opposite corners are empty (do not contain other terminals of X). Again the location of the points of X in all regions of the plane except the two opposite quadrants is uniquely determined. To compute the location of the remaining terminals in these two quadrants we adapt the second part of the algorithm of Edmonds [9]: we construct on these terminals a graph as in [9], partition it into connected components, separately determine the location of the points of each component, and then combine them into a single chain of components in order to obtain a global isometric embedding φ of (X, d) extending φ_0 or to decide that it does not exist.

Now, we briefly overview the algorithms of Edmonds [9] and Eppstein [10]. Edmonds [9] starts by picking two diametral points p, q of X . These two points can be embedded into the l_1 -plane in an infinite number of different ways. Each embedding defines an axis-parallel rectangle Π whose half-perimeter is exactly $d(p, q)$. Using the distances of p and q to the remaining points of X , Edmonds computes a list Δ of linear size of possible values of the sides of the rectangle Π . For each value δ from this list, the algorithm of [9] decides in $O(n^2)$ time if there exists an isometric embedding of X such that one side of the rectangle Π has length δ . For this, it partitions the points of X into groups, depending on their location in the regions of the plane determined by Π . In order to fix the positions of points in one of these regions, Edmonds [9] defines a graph whose connected components are also used in our algorithm. While sweeping through the list Δ , the algorithm of [9] updates this graph and its connected components in an efficient way. Notice that the second part of our algorithm is similar to that from [9], but instead of trying several sizes of the rectangle Π , we use the tight spans to provide us with a single rectangle, ensuring some rigidity in the embedding of the remaining points. The




FIGURE 3. The tree-network A_i .

algorithm of Eppstein [10] is quite different in spirit from our algorithm and that of Edmonds [9]. Eppstein [10] first incrementally constructs in $O(n^2)$ time a planar rectangular complex which is the tight span of the input metric space (X, d) or decide that the tight span of X is not planar. In the second stage of the algorithm, he decides in $O(n^2)$ time if this planar rectangular complex can be isometrically embedded into the l_1 -plane or not.

3.2. Computing the quadruplet P° . For each $i = 1, \dots, n$, set $X_i := \{x_1, \dots, x_i\}$. We start by computing the tight span of the first four points of X . If this tight span is not degenerated then we return the quadruplet X_4 as P° . Now suppose that the tight span of the first $i - 1$ points of X is a tree-network A_{i-1} with at most four leaves. This means that A_{i-1} contains one or two ramification points (which are not necessarily points of X) having degree at most 4, all remaining terminals of X_{i-1} are either leaves or vertices of degree two of A_{i-1} . We say that two terminals of X_{i-1} are consecutive in A_{i-1} if the segment connecting them in A_{i-1} does not contain other points of X_{i-1} . Note that A_{i-1} contains at most $n + 4$ of consecutive pairs. For each pair x_j, x_k of consecutive terminals of X_{i-1} we compute the Gromov product $\alpha_{x_i} := (x_j, x_k)_{x_i} = 1/2(d(x_i, x_j) + d(x_i, x_k) - d(x_j, x_k))$ of x_i with $\{x_j, x_k\}$. Let $\{a, b\}$ be the pair of consecutive points of X_{i-1} minimizing the Gromov product $\alpha_{x_i} = (a, b)_{x_i}$. Let c be the point of the segment $[a, b]$ of A_{i-1} located at distance $\alpha_a := (b, x_i)_a$ from a and at distance $\alpha_b := (a, x_i)_b$ from b (c may coincide with one of the points a or b).

Denote by A_i the tree-network obtained from A_{i-1} by adding the segment $[x_i, c]$ of length α_{x_i} . By running Breadth-First-Search on A_i rooted at x_i , we check if $d_{A_i}(x_i, x_j) = d(x_i, x_j)$ for any terminal x_j of X_i . If this holds for all $x_j \in X_i$, then the tight span of X_i is the tree-network A_i . If A_i contains more than 4 leaves, then we return the answer “not” and the algorithm halts. Otherwise, if $i = n$, then we return the answer “yes” and an isometric embedding of X and its tight span A_n in the l_1 -plane, else, if $i < n$, we consider the next point x_{i+1} . Finally, if x_j is the first point of X_i such that $d_{A_i}(x_i, x_j) \neq d(x_i, x_j)$, then we assert that *the tight span of the quadruplet $\{a, b, x_i, x_j\}$ is non-degenerated and we return it as P°* . Suppose by way of contradiction that $T(P^\circ)$ is a tree. Since A_{i-1} realizes X_{i-1} and $T(P^\circ)$ realizes P° , the subtree of A_{i-1} spanned by the terminals a, b , and x_j is isometric to the subtree of $T(P^\circ)$ spanned by the same terminals. On the other hand, $T(P^\circ)$ contains a point c' located at distance α_{x_i} , α_a , and α_b from x_i , a , and b , respectively. This means that $T(P^\circ)$ is isometric to the subtree of A_i spanned by the vertices x_i, a, b , and x_j , (see Fig. 3) contrary to the assumption that $d_{A_i}(x_i, x_j) \neq d(x_i, x_j)$. Hence, this inequality implies indeed that $T(P^\circ)$ is not a tree. Finally note that dealing with a current point x_i takes time linear in i , thus the whole algorithm for computing the quadruplet P° runs in $O(n^2)$ time.

3.3. Classification of the points of X with respect to the rectangle of $T(P^\circ)$. Let $P^\circ = \{p_1^\circ, p_2^\circ, p_3^\circ, p_4^\circ\}$ be the quadruplet whose tight span $T(P^\circ)$ is non-degenerated. Let R° be one of the two possible isometric embeddings of the rectangle $R(P^\circ)$ of $T(P^\circ)$ and consider a complete or a partial isometric embedding of $T(P^\circ)$ such that $R(P^\circ)$ is embedded as R° . Denote by $Q_1^\circ, Q_2^\circ, Q_3^\circ, Q_4^\circ$ the four (closed) quadrants defined by the four consecutive corners $q_1^\circ, q_2^\circ, q_3^\circ, q_4^\circ$ of R° labeled in such a way that the point p_i° must be located in the quadrant $Q_i^\circ, i = 1, \dots, 4$. Let also $S_1^\circ, S_2^\circ, S_3^\circ$, and S_4° be the remaining half-infinite strips. Since we know how to construct in constant time the tight span of a 5-point metric space, we can compute the distances from all terminals p of X to the corners of the rectangle $R(P^\circ)$ (and hence to the corners of R°) in total $O(n)$ time. With some abuse of notation, we will denote the l_1 -distance from p to the corner q_i° of R° by $d(p, q_i^\circ)$. Since R° is gated, from the distances of p to the corners of R° we can compute the gate of p in R° . Consequently, for each point $p \in X \setminus P^\circ$ we can decide in which of the nine regions of the plane will belong its location $\varphi(p)$ under any isometric embedding φ of (X, d) subject to the assumption that $R(P^\circ)$ is embedded as R° . If $\varphi(p)$ belongs to one of the four half-strips or to R° , then we can also easily find the exact location itself: this can be done by using either the gate of p in R° or the fact that inside these five regions the intersection of the four l_1 -spheres centered at the corners of R° and having the distances from respective corners to p as radii is a single point. So, it remains to decide the locations of points assigned to the four quadrants $Q_1^\circ, Q_2^\circ, Q_3^\circ$, and Q_4° . For any point $p \in X$ which must be located in the quadrant Q_i° , the set of possible locations of p is either empty (and no isometric embedding exists) or a segment s_p of Q_i° consisting of all points $z \in Q_i^\circ$ such that $\|z - q_i^\circ\|_1 = d(p, q_i^\circ)$.

Notice that for any quadruplet $P' = \{p'_1, p'_2, p'_3, p'_4\}$ of terminals such that p'_i is assigned to the quadrant $Q_i^\circ, i \in \{1, 2, 3, 4\}$, the rectangle R° belongs to the tight span $T(P')$ of P' . Indeed, for any point $p'_i, i \in \{1, 2, 3, 4\}$ and any point r of R° , we have $\|p'_i - r\|_1 + \|r - p'_j\|_1 = \|p'_i - p'_j\|_1$, where j is selected in such a way that q_i° and q_j° are opposite corners of R° . From injectivity of the l_1 -plane and the characterization of tight spans we conclude that all points of R° belong to $T(P')$, establishing in particular that this tight span is also non-degenerated.

3.4. The quadruplet P and its properties. Let $P = \{p_1, p_2, p_3, p_4\}$ be the quadruplet of X , where p_i is a point of X which must be located in the quadrant Q_i° and is maximally distant from the corner q_i° of R° . As we established above, the tight span of P is non-degenerated, moreover the rectangle $R(P)$ contains the rectangle $R(P^\circ)$. As we also noticed, there exists a constant number of ways in which we can isometrically embed $T(P)$ into the l_1 -plane. Further we proceed in the following way: we pick an arbitrary isometric embedding φ_0 of $T(P)$ and try to extend it to an isometric embedding φ of the whole metric space (X, d) in the l_1 -plane. If this is possible for some embedding of $T(P)$, then the algorithm returns the answer “yes” and an isometric embedding of X , otherwise the algorithm returns the answer “not”. Let R denote the image of $R(P)$ under φ_0 .

We call a terminal p_i of P *fixed* by the embedding φ_0 if either $\varphi_0(p_i)$ is a corner of the rectangle R or the segment of $T(P)$ incident to p_i is embedded by φ_0 as a horizontal or a vertical segment; else we call p_i *free*. The embedding of a free terminal p_i is not exactly determined but is restricted to a segment s_{p_i} consisting of the points of the quadrant defined by q_i and having the same l_1 -distance to q_i . We call the terminals $p_i, p_{i+1(\text{mod } 4)}$ *incident* and the terminals $p_i, p_{i+2(\text{mod } 4)}$ *opposite*. From the isometric embedding of $T(P)$ we conclude that




FIGURE 4. Possible isometric embeddings of $T(P)$.




FIGURE 5. The partition of the plane into half-strips and quadrants.

at most one of two incident terminals can be free. Moreover, if a terminal p_i of P is fixed but is not a corner of R , then at least one of the two terminals incident to p_i is also fixed. If all four tips of $T(P)$ are non-degenerated, then all four terminals of P are fixed. If only three tips of $T(P)$ are non-degenerated then at most one terminal of P is free, all remaining terminals are fixed. If only two tips of $T(P)$ are non-degenerated, then either they correspond to incident terminals, one of which is fixed and another one is free or to two opposite terminals which are both free. Finally, if only one tip of the tight span is non-degenerated, then it corresponds to a free terminal, all other terminals of P are corners of R and therefore are fixed (see Fig. 4 for the occurring possibilities).

Denote by Π the smallest axis-parallel rectangle containing R and the fixed terminals of P ; Fig. 5 illustrates Π for two cases from Fig. 4 (if a terminal is free, then the respective corner of R is also a corner of Π). Let q_1, q_2, q_3, q_4 be the corners of Π labeled in such a way that q_i is the corner of R corresponding to the point p_i and to the corner q_i° of R° . Denote by Q_1, \dots, Q_4 the quadrants of \mathbb{R}^2 defined by the corners of Π and by S_1, \dots, S_4 the remaining half-infinite strips. Again, as in the case of the quadruplet P° , by building the tight spans of $P \cup \{p\}$ for all terminals $p \in X \setminus P$, we can compute in total linear time the distances from all such points p to the corners of R (and to the corners of Π). From these four distances and the distances of p to the terminals of the quadruplet P we can determine in which of the nine regions $Q_1, Q_2, Q_3, Q_4, S_1, S_2, S_3, S_4, \Pi$ of the plane must be located p . Moreover, if p is assigned to the rectangle Π or to one of the four half-strips S_1, S_2, S_3, S_4 , then we can conclude




FIGURE 6. On possible locations of terminals in Q_1 and Q_3 .

that, in the region in which p assigned, the intersection of the four spheres centered at the terminals of P and having the distances from respective points to p as radii is either empty or a single point. The sphere centered at a free terminal p_i is needed only to decide the location of p in the quadrant Q of the plane having the same apex a as the quadrant Q_i and which is opposite to Q_i (a is a corner of Π). But in this case, instead of considering the sphere of radius $d(p, p_i)$ centered at $\varphi_0(p_i)$ we consider the sphere of radius $d(p, p_i) - \|\varphi_0(p_i) - a\|_1$ and centered at a : indeed, both these spheres have the same intersection with Q .

We are now ready to prove the following property of the quadruplet P : *among the four quadrants Q_1, Q_2, Q_3 , and Q_4 defined by P , two opposite quadrants are empty*, i.e., they do not contain terminals of $X \setminus P$. First note that by inspecting the different cases listed in Fig. 4 one can check that the two neighbors $p_{i-1(\bmod 4)}$ and $p_{i+1(\bmod 4)}$ of a free point $p_i \in P$ are both fixed; let say p_1 and p_3 are fixed. Now, suppose by way of contradiction that a terminal $q \in X \setminus P$ must be located in the quadrant Q_1 . This means that its gate in the rectangle Π is the corner of Π corresponding to p_1 . Since in any embedding φ of X that extends the chosen embedding of $T(P^\circ)$ the terminal p_1 is located in Q_1° , we deduce that $Q_1(\varphi(p_1)) \subseteq Q_1^\circ$. On the other hand, the inclusion $Q_1 \subseteq Q_1(\varphi(p_1))$ follows directly from the definition of Q_1 and the fact that p_1 is fixed. Now, from the inclusions $Q_1 \subseteq Q_1(\varphi(p_1)) \subseteq Q_1^\circ$, we obtain that $q \in Q_1^\circ$ and, since q is closer to p_1 than to q_1° , we get a contradiction with the choice of p_1 , establishing that indeed Q_1 does not contain any point of $X \setminus P$. The same argument shows that Q_3 is empty as well. Note that actually we proved that any quadrant Q_i corresponding to a fixed terminal p_i of P is empty.

3.5. Locating in the non-empty quadrants Q_1 and Q_3 . As we have showed in previous subsection, any isometric embedding φ of (X, d) extending the embedding φ_0 of $T(P)$ locates each terminal p of $X \setminus P$ in one and the same of the nine regions defined by Π . Moreover, if p must be located in the rectangle Π or in one of the four half-strips S_1, \dots, S_4 , then this location $\varphi(p)$ is uniquely determined from the distances to the terminals of P and to the corners of Π . We also established that only one or two opposite quadrants defined by Π , say Q_1 and Q_3 , can host terminals of $X \setminus P$; see Fig. 6. We will show now how to find the exact location of the set X_1 of terminals assigned to Q_1 (the set X_3 of terminals which must be located in Q_3 is treated analogously).

Note that independently of how the extension φ of φ_0 is chosen, for each terminal $u \in X_1$, the l_1 -distance $\|\varphi(u) - q_1\|_1$ from the location of u to the corner q_1 of Π is one and the same,




FIGURE 7. $\varphi(u)$ and $\varphi(v)$ are fixed by $\varphi(p^*)$ and $\varphi(p_*)$




FIGURE 8. $\varphi(v)$ is fixed by $\varphi(p^*)$

which we denote by Δ_u . The value of Δ_u can be easily computed because q_1 lies between $\varphi(u)$ and $\varphi(p_i)$ for any $p_i \in P$: for example, we can set $\Delta_u := d(u, p_1) - \|\varphi_0(p_1) - q_1\|_1$. Then the set of all possible locations $\varphi(u)$ of $u \in X_1$ is the *level segment* s_u which is the intersection of Q_1 with the sphere $S(q_1, \Delta_u)$ of radius Δ_u centered at q_1 .

To compute the locations of the terminals of X_1 in the quadrant Q_1 , we adapt to the l_1 -plane the definition of a graph (which we denote by $G_1 = (X_1, E_1)$) defined by Edmonds [9] in the l_∞ -plane. Two terminals $u, v \in X_1$ are adjacent in G_1 if and only if $d(u, v) > |\Delta_u - \Delta_v|$. Equivalently $u, v \in X_1$ with $\Delta_u \leq \Delta_v$ are adjacent in G_1 iff u cannot be located between q_1 and v : $\varphi(u) \notin I_1(q_1, \varphi(v))$. Denote by C_1, C_2, \dots, C_k the connected components of the graph G_1 . They have the following useful properties established in Lemmata 3-5 of [9]:

- (1) Each component C_i is *rigid*, i.e., once the location of any point u of C_i has been fixed, the locations of the remaining points of C_i are also fixed (up to symmetry with respect to the line parallel to the bisector of Q_1 and passing via u);
- (2) The components C_1, C_2, \dots, C_k of the graph G_1 can be numbered so that the points of each C_i appear consecutively in the list of points $u \in X_1$ sorted in increasing order of their distances Δ_u to q_1 ;
- (3) For a component C_i of G_1 , let B_i be the smallest axis-parallel rectangle containing $\{\varphi_i(u) : u \in C_i\}$ for an isometric embedding φ_i of (C_i, d) in the l_1 -plane. Let b_i be the upper right corner of B_i . Then the embedding of C_1, C_2, \dots, C_k preserves the distances between all pairs of points lying in different components if and only if for every pair of consecutive components C_i and C_{i+1} , the rectangle B_{i+1} lies entirely in the quadrant $Q_1(b_i)$.

The location in the quadrant Q_1 of some terminals of X_1 (and therefore of the connected components containing them) can be fixed by terminals already located in the two half-strips incident to Q_1 . We say that a terminal $u \in X_1$ is *fixed by a terminal* p already located in $S_1 \cup S_4$ if the intersection of the segment s_u with the sphere $S(\varphi(p), d(p, u))$ is a single point. Note that if $u \in X_1$ is fixed by a terminal located in S_1 , then u is also fixed by the upmost terminal p^* located in this half-strip. Analogously, if $u \in X_1$ is fixed by a terminal of S_4 , then u is also fixed by the rightmost terminal p_* located in S_4 . Therefore by considering the intersections of the segments $s_u, u \in X_1$, with the spheres $S(\varphi(p^*), d(p^*, u))$ and $S(\varphi(p_*), d(p_*, u))$ we can decide in linear time which terminals of X_1 are fixed by p^* and p_* and find their location in Q_1 (for an illustration, see Fig. 7). According to property (1), if a terminal of a connected component




FIGURE 9. On the assemblage of blocks B_{j+1}, \dots, B_k .

of G_1 is fixed, then the location of the whole component is also fixed (up to symmetry). Let C_j be the connected component of G_1 containing the furthest from q_1 terminal $u \in X_1$ fixed by p^* or p_* , say by p^* (therefore the location of C_j is fixed). We assert that *all terminals of C_1, C_2, \dots, C_{j-1} are also fixed by p^** . Indeed, pick such a terminal v . From property (2) we conclude that $\Delta_v \leq \Delta_u$ and from the definition of G_1 we deduce that v must be located in the axis-parallel rectangle $I_1(q_1, \varphi(u))$, and therefore below u . Since u is fixed by p^* , u must be located below p^* , whence v also must be located below p^* . We can easily see that the intersection of s_v with the sphere $S(\varphi(p^*), d(p^*, v))$ is a single point, i.e. v is also fixed by p^* (see Fig. 8).

It remains to locate in Q_1 the terminals of the components $C_{j+1}, C_{j+2}, \dots, C_k$. We compute separately an isometric embedding of each component C_i for $i = j+1, \dots, k$. For this, we fix arbitrarily the location of the first two points u, v of C_i in the segments s_u and s_v so that to preserve the distance $d(u, v)$ (the terminals of C_i are ordered by their distances to q_1). By property (1) of [9], the location of the remaining points of C_i is uniquely determined and each point w of C_i will be located in its level segment s_w . Let φ_i be the resulting embedding of C_i . Denote by B_i the smallest axis-parallel rectangle (alias *box*) containing the image $\varphi_i(C_i)$ of C_i . Let a_i and b_i denote the lower left and the upper right corners of B_i . Note that a_i belongs to the l_1 -interval between q_1 and the image $\varphi_i(u)$ of any terminal u of C_i , while the l_1 -interval between q_1 and b_i will contain the images of all terminals of C_i . Therefore if we set $\Delta_{a_i} := \Delta_u - \|a_i - \varphi_i(u)\|_1$ and $\Delta_{b_i} := \Delta_u + \|\varphi(u) - b_i\|_1$, where u is any terminal of C_i , then in all isometric embeddings of (C_i, d) in which all terminals $u \in C_i$ are located on s_u , the points a_i and b_i must be located on the level segments s_{a_i} and s_{b_i} , defined as the intersections of the quadrant Q_1 with the spheres $S(q_1, \Delta_{a_i})$ and $S(q_1, \Delta_{b_i})$.

By properties (2) and (3) of [9], in order to define a single isometric embedding of the components C_{j+1}, \dots, C_k we now need to assemble the boxes B_{j+1}, \dots, B_k (by moving their terminals along the level segments) in such a way that *for two consecutive components C_i and C_{i+1} , the box B_{i+1} lies entirely in the quadrant $Q_1(b_i)$* . We assert that this is possible if and only if *for each pair of consecutive boxes B_i, B_{i+1} , $i = j, j+1, \dots, k-1$, the inequality $\Delta_{b_i} \leq \Delta_{a_{i+1}}$ holds*. Indeed, if $\Delta_{b_i} \leq \Delta_{a_{i+1}}$, then translating B_{i+1} along the segment $s_{a_{i+1}}$, we can locate its corner a_{i+1} in the quadrant $Q_1(b_i)$ and thus satisfy the embedding requirement. Conversely, if $\Delta_{b_i} > \Delta_{a_{i+1}}$ holds, then a_{i+1} cannot belong to the quadrant $Q_1(b_i)$ independently of the

positions of a_{i+1} and b_i on their level segments. This local condition depends only of the values of $\Delta_{a_i}, \Delta_{b_i}$ and is independent of the actual location of the boxes $B_i, i = 1, \dots, k$. As a result, the algorithm that embeds the boxes B_{j+1}, \dots, B_k is very simple. For each $i = j, \dots, k-1$, we compute the box B_{i+1} and the values of $\Delta_{a_{i+1}}$ and $\Delta_{b_{i+1}}$. If $\Delta_{a_{i+1}} < \Delta_{b_i}$ for some i , then return the answer “there is no isometric embedding of (X, d) extending the embedding φ_0 of $T(P)$ ”. Otherwise, having already located the box B_i , by what has been shown above, the intersection of the quadrant $Q_1(b_i)$ with the level segment $s_{a_{i+1}}$ is non-empty. Therefore we can translate B_{j+1} in such a way that its lower left corner a_{i+1} becomes a point of this intersection.

In this way, we obtain an embedding of C_{j+1}, \dots, C_k and B_{j+1}, \dots, B_k satisfying the conditions (1)-(3), thus an isometric embedding of the metric space $(\bigcup_{i=j+1}^k C_i, d)$ in Q_1 . Analogously, by constructing the graph $G_3 = (X_3, E_3)$ and its components, either we obtain a negative answer or we return an isometric embedding of the metric space defined by the non-fixed components of G_3 in the quadrant Q_3 . Denote by φ the embedding of X which coincides with φ_0 on P , with these two embeddings on the non-fixed components of G_1 and G_3 , and with the already computed fixed locations of the terminals assigned to Π , to the half-strips S_1, S_2, S_3, S_4 , and to the fixed connected components of the graphs G_1 and G_3 . In $O(n^2)$ we test if φ is an isometric embedding of (X, d) into the l_1 -plane. If the answer is negative, then we return “there is no isometric embedding of (X, d) extending the embedding φ_0 of $T(P)$ ”, otherwise we return φ as an isometric embedding. The algorithm returns the global answer “not” if for all possible embeddings φ_0 of $T(P)$ it returns the negative answer. From what we established follows that in this case (X, d) is not isometrically embeddable into the l_1 -plane.

3.6. Algorithm and its complexity. We conclude the paper with a description of the main steps of the algorithm and their complexity.

Algorithm Embedding into the l_1 -plane

Input: A metric space (X, d) on n points

Output: An isometric embedding φ of (X, d) into (\mathbb{R}^2, d_1) or the answer “not” if it does not exist

Step 1. Find a quadruplet P° of X whose tight span contains a rectangle. If P° does not exist, then $T(X)$ is a tree. If $T(X)$ has more than 4 leaves, then return “not”, else return an embedding of $T(X)$ and (X, d) .

Step 2. Pick any embedding of $T(P^\circ)$ and for each terminal of $X \setminus P^\circ$ determine in which of the nine regions of the plane it must be located. Using this partition of $X \setminus P^\circ$, define the quadruplet P .

Step 3. Embed P and its tight span $T(P)$ into the l_1 -plane in all possible different ways. Try to extend each of these embeddings to an isometric embedding of (X, d) following the rules (a)-(g). If all of these attempts return the answer “not”, then return the answer “not”, else return one of the obtained embeddings.

- (a) Given an embedding φ_0 of $T(P)$, for each terminal u of $X \setminus P$ determine in which of the nine regions defined by the rectangle Π will be located u in any isometric embedding extending φ_0 ;
- (b) Locate the terminals assigned to the rectangle Π and the four half-strips S_1, S_2, S_3, S_4 ;
- (c) Define the sets of terminals X_1 and X_3 assigned to the quadrants Q_1 and Q_3 , construct the graphs $G_1 = (X_1, E_1)$ and $G_3 = (X_3, E_3)$ and their connected components;
- (d) Find the terminals of X_1 fixed by p^*, p_* and their location in Q_1 . Do a similar thing for X_3 ;
- (e) Find an isometric embedding of each component C_i of G_1 not containing already fixed terminals so that its terminals are located on their level segments. Do a similar thing for G_3 ;
- (f) Test if the free components C_{j+1}, \dots, C_k of G_1 satisfy the condition $\Delta_{b_i} \leq \Delta_{a_{i+1}}$ for $i = j+1, \dots, k-1$. If not, then return the answer “not”, else locate consecutively the boxes B_{j+1}, \dots, B_k in such a way that a_{i+1} is located in $Q_1(b_i) \cap s_{a_{i+1}}$ and fix in this way the position of all terminals of X_1 . Do a similar thing for the free components of G_3 ;
- (g) Verify if the resulting embedding of X extending φ_0 is an isometric embedding of (X, d) . If “yes”, then return it as a resulting isometric embedding, otherwise return the answer “there is no isometric embedding of (X, d) extending the embedding φ_0 ”.

In Subsection 3.2 we established that the quadruplet P° , if it exists, can be computed in $O(n^2)$ time. If P° does not exist, then the tree-network A_n (constructed within the same time bounds) is the tight span of (X, d) . Embedding A_n (if it has at most 4 leaves) in the l_1 -plane can be easily done in linear time. As shown in Subsection 3.3, Step 2 can be implemented in linear time. There exists a constant number of ways in which the quadruplet P and its tight span can be isometrically embedded in the l_1 -plane. Therefore, to show that Step 3 has complexity $O(n^2)$, it suffices to estimate the total complexity of the steps (a)-(g) for a fixed embedding φ_0 of $T(P)$. Step (a) is similar to Step 2, thus its complexity is linear. The exact location of each terminal in the half-strips or in Π is determined as the intersection of two spheres, therefore step (b) is also linear. Defining the graph G_1 and computing its connected components can be done in $O(|X_1|^2)$ time. Thus step (c) has complexity $O(n^2)$. Steps (d) and (e) can be implemented in an analogous way as (b), thus their complexity is $O(n)$. Testing the condition in step (f) and assembling the free components into a single chain is linear as well. Finally, step (g) requires $O(n^2)$ time. Therefore, the total complexity of the algorithm is $O(n^2)$. Summarizing, here is the main result of this note:

Theorem 1. *For a metric space (X, d) on n points, it is possible to decide in optimal $O(n^2)$ time if (X, d) is isometrically embeddable into the l_1 -plane and to find such an embedding if it exists.*

REFERENCES

- [1] N. Aronszajn and P. Panitchpakdi, Extensions of uniformly continuous transformations and hyperconvex metric spaces, *Pacific J. Math.* **6** (1956), 405–439.
- [2] D. Avis, M. Deza, The cut cone, L_1 embeddability, complexity and multicommodity flows, *Networks* **21** (1991), 595–617.
- [3] H.-J. Bandelt and V. Chepoi, Embedding metric spaces in the rectilinear plane: a six-point criterion, *Discr. Comput. Geom.* **15** (1996) 107–117.
- [4] H.-J. Bandelt, V. Chepoi, Embedding into the rectilinear grid, *Networks* **32** (1998), 127–132.
- [5] G.E. Christopher and M.A. Trick, Faster decomposition of totally decomposable metrics with applications, Preprint, Carnegie Mellon University (1996).
- [6] M. Deza, M. Laurent, *Geometry of Cuts and Metrics*, Springer-Verlag, Berlin, 1997.
- [7] A.W.M. Dress, Trees, tight extensions of metric spaces, and the cohomological dimension of certain groups, *Adv. Math.* **53** (1984), 321–402.
- [8] A. Dress and R. Scharlau, Gated sets in metric spaces, *Aequat. Math.* **34** (1987), 112–120.
- [9] J. Edmonds, Embedding into l_∞^2 is easy, embedding into l_∞^3 is NP-complete, *Discr. Comput. Geom.* **39** (2008), 747–765.
- [10] D. Eppstein, Optimally fast incremental Manhattan plane embedding and planar tight span construction, Electronic preprint arXiv:0909.1866v1, (2009).
- [11] J. Isbell, Six theorems about metric spaces, *Comment. Math. Helv.* **39** (1964), 65–74.
- [12] S.M. Malitz and J.I. Malitz, A bounded compactness theorem for L^1 -embeddability of metric spaces in the plane, *Discr. Comput. Geom.* **8** (1992) 373–385.